

WE CLAIM:

1. A method for controlling network congestion, the method comprising:
 - 5 receiving, at an edge device that is in communication with a Fibre Channel network fabric but is not part of the fabric, a frame having the following:
 - a source identifier corresponding to a destination node;
 - a destination identifier corresponding to the edge device;
 - an exchange identifier corresponding to an exchange being transmitted
 - 10 by the edge device to the destination node; and
 - an instruction from a node within the Fibre Channel network to the edge device, the instruction indicating that an exchange originated by the edge device is causing network congestion; and
 - 15 implementing a congestion reaction mechanism at the edge device in accordance with the instruction.
2. The method of claim 1, wherein the instruction pertains to the exchange being transmitted by the edge device to the destination node.
- 20 3. The method of claim 1, wherein the congestion reaction mechanism is applied to the exchange being transmitted by the edge device to the destination node.
4. A method for controlling network congestion, the method comprising:
 - 25 receiving, at a network device that is part of a Fibre Channel network device fabric, a frame having the following:
 - a source identifier corresponding to a source node, the source node comprising an edge device outside the Fibre Channel network device fabric;
 - a destination identifier corresponding to a destination node; and
 - 30 an exchange identifier corresponding to an exchange being transmitted by the source node to the destination node;
 - characterizing traffic flow at the Fibre Channel network device; and
 - sending an instruction from the Fibre Channel network device to the source node.

5. The method of claim 4, wherein the instruction sent to the source node comprises an Edge Quench message.

5 6. The method of claim 4, wherein the instruction has a source identifier field corresponding to the destination node, a destination identifier field corresponding to the source node, and an exchange identifier field corresponding to the exchange identifier field.

10 7. The method of claim 4, wherein the instruction contains network device congestion information.

8. The method of claim 4, wherein the instruction contains network device queue level information.

15 9. The method of claim 4, wherein the instruction directs the source node to reduce a rate of transmission.

20 10. The method of claim 9, wherein the exchange of which the rate of transmission is to be reduced is identified by an exchange identifier field of the instruction.

25 11. A method for controlling traffic flow between a first end node and a second end node through an intermediate node that is part of a Fibre Channel switch fabric, the first end node and the second end node being outside the Fibre Channel switch fabric, the method comprising:

transmitting a first frame having a source identifier corresponding to the first end node, a destination identifier corresponding to a second end node and a exchange identifier corresponding to a particular exchange being transmitted by the first end node to the second end node;

30 receiving a second frame from the intermediate node, the second frame having a source identifier corresponding to the second end node, a destination identifier corresponding to the first end node and an exchange identifier corresponding to the exchange identifier of the first frame, the second frame also including instructions to

adjust a current allowed rate for the exchange corresponding to the exchange identifier being transmitted from the first end node to the second end node; and

adjusting the current allowed rate for the particular exchange from the first end node to the second end node according to the instructions.

5

12. The method of claim 11, wherein the adjusting step is performed by the first end node.

10 13. The method of claim 11, wherein the current allowed rate cannot exceed a maximum allowed rate.

14. The method of claim 11, further comprising the step of determining that the second frame is an Edge Quench message.

15 15. The method of claim 11, wherein the adjusting step comprises decreasing the current allowed rate according to a decrease function.

16. The method of claim 11, wherein the adjusting step comprises increasing the current allowed rate according to an increase function.

20

17. The method of claim 16, wherein the increase function is dynamically set.

18. The method of claim 16, wherein the increase function is set based on a state of a queue associated with the intermediate node.

25

19. A device for controlling network congestion, the device comprising:
means for receiving, at an edge device that is in communication with a Fibre Channel network fabric but is not part of the fabric, a frame having the following:

a source identifier corresponding to a destination node;

30

a destination identifier corresponding to the edge device;

an exchange identifier corresponding to an exchange being transmitted by the edge device to the destination node; and

an instruction from a node within the Fibre Channel network device to the edge device, the instruction indicating that an exchange originated by the edge device is causing network congestion; and

means for implementing a congestion reaction mechanism at the edge device in accordance with the instruction.

20. A device for controlling network congestion, the device comprising:

means for receiving, at a network device that is part of a Fibre Channel network fabric, a frame having the following:

10 a source identifier corresponding to a source node, the source node comprising an edge device outside the Fibre Channel network device fabric;

 a destination identifier corresponding to a destination node; and

 an exchange identifier corresponding to an exchange being transmitted

15 by the source node to the destination node;

 means for characterizing traffic flow at the Fibre Channel network device; and

 means for sending an instruction from the Fibre Channel network device to the source node.

20 21. A device for controlling network congestion, the device comprising:

 a port for receiving, at a network device that is part of a Fibre Channel network fabric, a frame having the following:

 a source identifier corresponding to a source node, the source node comprising an edge device outside the Fibre Channel network device fabric;

25 a destination identifier corresponding to a destination node; and

 an exchange identifier corresponding to an exchange being transmitted

 by the source node to the destination node; and

 at least one processor configured to characterize traffic flow at the Fibre

30 Channel network device and to send an instruction from the Fibre Channel network device to the source node.

22. A host bus adaptor for controlling network congestion, the host bus adaptor comprising:

a first connection for communicating with a Fibre Channel network device and configured to receive frames comprising the following:

- 5 a source identifier corresponding to a destination node;
- a destination identifier corresponding to the edge device;
- an exchange identifier corresponding to an exchange being transmitted by the edge device to the destination node; and
- an instruction from a node within the Fibre Channel network device to the edge device, the instruction indicating that an exchange originated by the edge device is causing network congestion;
- 10 a second connection for communicating with a bus of the edge device; and
- at least one processor for implementing a congestion reaction mechanism at the edge device in accordance with the instruction.

23. A computer program embodied in a machine-readable medium, the computer
15 program including instructions for controlling one or more devices to alleviate network
 congestion by performing the following steps:

 receiving, at an edge device that is in communication with a Fibre Channel
network fabric but is not part of the fabric, a frame having the following:

- 20 a source identifier corresponding to a destination node;
- a destination identifier corresponding to the edge device;
- an exchange identifier corresponding to an exchange being transmitted by the edge device to the destination node; and
- an instruction from a node within the Fibre Channel network device to the edge device, the instruction indicating that an exchange originated by the edge device is causing network congestion; and
- 25 implementing a congestion reaction mechanism at the edge device in accordance with the instruction.

24. The computer program of claim 23, wherein the instruction pertains to the
30 exchange being transmitted by the edge device to the destination node.

25. The computer program of claim 23, wherein the congestion reaction mechanism
 is applied to the exchange being transmitted by the edge device to the destination node.

26. A computer program embodied in a machine-readable medium, the computer program including instructions for controlling one or more devices to alleviate network congestion by performing the following steps:

5 receiving, at a network device that is part of a Fibre Channel network fabric, a frame having the following:

a source identifier corresponding to a source node, the source node comprising an edge device outside the Fibre Channel network device fabric;

10 a destination identifier corresponding to a destination node; and
an exchange identifier corresponding to an exchange being transmitted by the source node to the destination node;

characterizing traffic flow at the Fibre Channel network device; and
sending an instruction from the Fibre Channel network device to the source

15 node.